

TRAINING COURSE

Apache Kafka[®] Administration by Confluent

Course Objectives

During this hands-on course you will learn how:

- Kafka and the Confluent platform work, and how their main subsystems interact
- To set up, manage, monitor and tune your cluster
- To use industry best practices developed by the world's foremost Apache Kafka experts

Hands-on Training

Throughout the course, hands-on exercises reinforce the topics being discussed.

Exercises include:

Using Kafka's command-line tools

- Automating configuration
- Using Kafka's administrative tools
- Tuning Producer and Consumer performance
- Securing the cluster
- Building data pipelines with Kafka Connect

Prerequisites

Attendees should have a working knowledge of the Kafka architecture, either through:

- Prior experience, or
- By completing training beforehand to ensure familiarity with the relevant concepts. Visit www.confluent.io/training to learn the fundamentals of data streaming and Apache Kafka.

It is also important to have strong knowledge of Linux/Unix and understand basic TCP/IP networking concepts. Familiarity with Java Virtual Machine (JVM) is helpful.

Participants are required to provide a laptop computer with unobstructed internet access to fully participate in the class.

To evaluate your Kafka knowledge for this course, please complete the self assessment: <https://cnfl.io/fundamentals-quiz>

To sign-up for one of our courses, visit us [here](#).

Who Should Attend?

This course is designed for engineers, system administrators, and operations staff responsible for building, managing, monitoring, and tuning Kafka clusters.

Module	Description
Fundamentals of Apache Kafka	<ul style="list-style-type: none"> • Kafka as a Distributed Streaming Platform • The Distributed Log • Producer and Consumer Basics
Apache Kafka Architecture	<ul style="list-style-type: none"> • Kafka's Commit Log • Replication for High Availability • Partitions and Consumer Groups for Scalability • Security Overview
Providing Durability	<ul style="list-style-type: none"> • Data Replication • Failure Recovery • Log Files & Offset Management • Exactly-Once Semantics (EOS)
Managing a Cluster	<ul style="list-style-type: none"> • Installing and Running Kafka • Configuration Management • Monitoring • Log Retention and Compaction • Commissioning and Decommissioning Brokers
Optimizing Apache Kafka Performance	<ul style="list-style-type: none"> • Monitoring, Testing, and Tuning Brokers and Kafka Clients • The Consumer Group Protocol
Apache Kafka Security	<ul style="list-style-type: none"> • Transport Encryption • Authentication • Securing Apache Kafka • Migrating to a Secure Cluster
Data Pipelines with Kafka Connect	<ul style="list-style-type: none"> • The Motivation for Kafka Connect • Types of Connectors • Kafka Connect Implementation • Standalone and Distributed Modes • Configuring the Connectors
Kafka in Production	<ul style="list-style-type: none"> • Kafka Reference Architecture for Apache Kafka and the Complete Confluent Platform • Capacity Planning • Multi Data Center Deployments

Visit confluent.io/training for more information.